

CAUSES, COUNTERFACTUALS, AND NON-LOCALITY

Mathias Frisch

In order to motivate the thesis that there is no single concept of causation that can do justice to all of our core intuitions concerning that concept, Ned Hall has argued that there is a conflict between a counterfactual criterion of causation and the condition of causal locality. In this paper I critically examine Hall's argument within the context of a more general discussion of the role of locality constraints in a causal conception of the world. I present two strategies that defenders of counterfactual accounts of causation can pursue to respond to Hall's challenge—including the adoption of a counterfactual condition that is sufficient for causal action-at-a-distance in place of Hall's 'process' condition—and conclude that Hall's argument against counterfactual accounts of causation is unsuccessful.

1. Introduction

What is it for one event to be a cause of another? It has become increasingly popular to argue that no single account of the concept of cause can adequately capture all our core intuitions concerning causation. The only way fully to do justice to our intuitions, it is suggested, is to allow for causal pluralism and to maintain that we employ several distinct causal concepts. Causal pluralism recommends itself, it is claimed, even if we restrict our attention to token-level causal claims. In this paper I want to focus on one influential argument in support of causal pluralism—Ned Hall's claim that there is a conflict between the condition that there is no action-at-a-distance and the idea that effects counterfactually depend on their causes, which he defends in [Hall 2002] and in the widely cited [Hall 2004]. According to Hall, counterfactual accounts of causation cannot distinguish between action-at-a-distance of a genuine and potentially problematic kind, on the one hand, and what he calls 'ersatz action-at-a-distance', on the other. If causation consisted in counterfactual dependence, then there would exist perfectly 'ordinary' circumstances that involved action-at-a-distance. Thus, there appears to be a tension between the claim that causation can be understood counterfactually and the desideratum that there be no genuine action-at-a-distance, since counterfactual accounts of causation

do not allow us to formulate a prohibition against action-at-a-distance in a way that does not exclude many intuitively causally unproblematic circumstances as well. Hall concludes from this apparent conflict that we have to distinguish at least two types of causal relation: Causation as counterfactual dependence (possibly spelled out along Lewisian lines), which plays no role in considerations of locality, and a distinct type of causal relation, causation as ‘production’, which is invoked in the requirement that causes do not act at a distance.

Worries about action-at-a-distance have a long tradition in physics and philosophy,¹ yet it is often unclear what motivates such worries and even what the worries’ precise content is meant to be. That is, neither is it made explicit what exactly it would be for a world to contain genuine action-at-a-distance nor are the reasons spelled out as to why we should find a world with such features objectionable. In particular, *causal* conditions of locality are often not carefully distinguished from *nomological* conditions and the precise relations between the two kinds of condition are left obscure. One of my aims in this paper is to make some progress toward a better understanding of the place of locality constraints in a causal conception of the world by carefully distinguishing several locality conditions, examining their relationships and showing how different accounts of causation suggest different ways of imposing locality constraints.

Hall’s thesis that counterfactual accounts of causation do not allow us to draw what appears to be an intuitively obvious distinction between local and non-local situations exposes a cluster of puzzles concerning locality constraints, including the following: are counterfactual accounts of causation in fact incapable of incorporating genuine locality constraints, as Hall argues? If indeed there are intuitively local situations that come out as non-local according to certain accounts of causation, then what does distinguish genuine action-at-a-distance from that of the ‘ersatz’ variety? Are there good reasons for demanding that an adequate account of causation should be able to incorporate such constraints? In what follows I will examine the first two questions in detail, but I also want to briefly address the third question.

¹ Recall, for example, Newton’s well-known reservations about his own theory of gravitational attraction: That one body may act upon another at a distance through a vacuum without the mediation of anything else [...] is to me so great an absurdity that, I believe, no man who has in philosophic matters a competent faculty of thinking could ever fall into it. (In a letter to Richard Bentley, dated February 25 1693, reprinted in [I.B. Cohen 1978: 302-3]).

It is not immediately obvious that the answer to this last question is ‘yes’. A defender of a counterfactual account might, for example, insist that the relevant distinction between local and non-local situations that is at the heart of our intuitive ‘distaste’ for action-at-a-distance is a purely *nomological* distinction and that our theory of *causation* need not be able to mirror that distinction. That is, one might reply that our intuitions regarding action-at-a-distance should properly be understood as being concerned with the character of the *laws* governing a particular world and that it is no constraint on an adequate theory of *causation* that it be able to provide an analogous distinction. Alternatively, one might maintain that the intuitive distinction between genuine and ersatz non-locality *is* causal and not (merely) nomological, but that it derives its plausibility entirely from an implicit commitment to a ‘production’ account of causation. But then it is open to defenders of counterfactual accounts to reject the relevant intuitions along with the latter account of causation. What these brief considerations suggest is that a defender of a counterfactual account can accept much of Hall’s argument without having to give up her account. She might concede that the intuitive distinction between local and non-local circumstances captures something of philosophically relevance but deny that the distinction is causal; or, alternatively, she could concede that the distinction is causal but reject it as being rooted in an unacceptable account of causation. Neither of these two responses commit counterfactual accounts to being able to reproduce a putatively intuitive distinction between local and genuinely non-local situations.

Thus, even if Hall were able to show that counterfactual accounts cannot not distinguish between genuine and ‘ersatz’ action-at-a-distance, this would not yet be enough to show that counterfactual accounts need to be abandoned. My main goal in this paper, however, is to argue that there is a third response on behalf of defenders of a counterfactual account: contrary to what Hall claims, counterfactual accounts can draw a distinction between local and non-local situations in intuitively the right place. Counterfactual accounts of causations may face a number of difficulties—and Hall himself discusses several—but being able to distinguish locally acting causes from those acting at a distance is not among them. Indeed, counterfactual accounts can employ several strategies for drawing the relevant distinction. One strategy, I will argue, appeals to a ‘field-theoretic’ account of property instantiation. A second strategy, and the one I prefer, involves rejecting Hall’s proposal

for a sufficient condition for action-at-a-distance. In place of Hall's condition I will propose a sufficient condition for action-at-a-distance that advocates of counterfactual accounts of causation can readily accept.

I will proceed as follows. In the next section I will distinguish different causal and nomological locality conditions and will summarize Hall's argument, which centrally involves an appeal to a certain causal structure—so called *double preventions*. In section 3 I will show that Hall's worry arises for causal structures even simpler than double preventions and will suggest a first reply to Hall's challenge based on a 'field-theoretic' way of construing talk of omissions. In section 4 I will present a counterfactual criterion for causal action-at-a-distance and defend the criterion against an objection suggested by a brief remark in [Hall 2002]. I will end with a brief conclusion.

2. A conflict between locality and counterfactual dependence?

Hall proposes the following sufficient condition (*SC*) for action-at-a-distance: 'We have a case of it if we have a cause, at least one of whose effects is not connected to it via any spatiotemporal continuous causal chain.' [Hall 2002: 276] He calls the absence of action-at-a-distance *Locality*. The second causal condition with which he is concerned is *Dependence*—a condition which, he believes, any counterfactual account of causation would want to endorse and which reads as follows: 'If for two actual events *c* and *e*, *e* counterfactually depends on *c*, then *c* is a cause of *e*.' Hall argues that there is a certain class of cases—cases of *double prevention*—that involve counterfactual dependence but violate *Locality*. The significance of this result, according to Hall, is that the cases he presents do not involve physical processes that we would intuitively take as involving action-at-a-distance. If we accept *Dependence* and, therefore, if we accept that the cases discussed by Hall are genuine instances of causation, then causal non-locality seems to be altogether too easily achievable—even in cases in which the physics in some intuitive sense appears to be perfectly local. If we accept *Dependence*, then we cannot, it seems, distinguish between genuine cases of action-at-a-distance and 'ersatz' cases, such as the ones discussed by Hall. That is, the problem Hall sees for counterfactual accounts is not simply that there are relations between events that are causal according to *Dependence* yet violate *Locality*—Hall does not think that action-at-a-distance is conceptually or metaphysically

impossible—but rather that there are situations that are local by some *other* standard and nevertheless, given *Dependence*, violate *Locality*.

How do we distinguish cases of genuine action-at-a-distance from those that involve only what Hall calls ‘ersatz’ action-at-a-distance? Curiously, Hall never says explicitly how he intends to draw this distinction. When he introduces the problem, he points to what he takes to be a paradigm example of genuine action-at-a-distance: a spell cast in a world where the ‘Laws of Magic’ determine when and where the spell will take effect without the mediation of a spatiotemporally continuous causal chain [Hall 2002: 277]. And toward the end of that paper [2002: 291] Hall considers a criterion of nomological locality, which he attributes to Adam Elga, that seems to be violated in the magic world. Elga’s condition of nomological locality presupposes determinism and states:

(NL_1) The physical laws are local exactly if the state of the world at a spacetime point P is nomologically determined by the state of the world on any arbitrarily small spatiotemporal shell S surrounding P .

Alternatively we can define nomological locality as follows [Belot 1998: 540]:

(NL_2) The laws governing a situation are local exactly if in order to predict what will happen *here* in a finite amount of time, Δt , we need only look at the present state of the world in a finite neighborhood of *here*, and the size of the neighborhood shrinks to zero as $\Delta t \rightarrow 0$.

NL_2 is strictly stronger than NL_1 , in that the second condition disallows dependences that propagate infinitely fast, but we can ignore the difference between the two conditions in what follows.²

Hall ultimately rejects Elga’s criterion as a substitute for his own condition SC since, as he argues, violations of the criterion only imply that action-at-a-distance is *permitted* but do not imply that action-at-a-distance actually *occurs*. But he also claims that when NL_1 is satisfied, then there is no causal action-at-a-distance. If NL_1 is satisfied, then, Hall says, ‘it follows that no event outside of [a spacetime shell S surrounding event e at P] acts directly at a distance on e .’ [Hall 2002:291]. Thus,

² See [Lange 2002] and [Frisch 2005] for discussions of various different locality conditions and their relations.

while Hall does not believe that NL_1 provides a necessary condition for genuine causal locality, he does seem to believe that it provides a *sufficient* condition. Instead, then, of having to rely merely on an appeal to a vague and intuitive distinction between genuine and ‘ersatz’ cases of action-at-a-distance, Hall might wish to argue that there can be no genuine causal action-at-a-distance in situations governed by laws that are local, according to either NL_1 or NL_2 .

It is far from clear, however, that Hall’s claim that nomological locality implies the absence of causal action-at-a-distance is correct. For example, it seems that we can get a failure of *Locality* even in a world with local laws, if there are events c in such a world with the following properties: c acts at distance on an event e , but c also has effects on at least some events in any spatiotemporal neighborhood of e . If c is a common cause both of e and of events in any neighborhood of e and if the occurrence of c is encoded in the state of any shell surrounding e in such a way that the state of an arbitrary shell is nomologically sufficient for the occurrence of e , then the system is nomologically local even though c acts on e at a distance.³ Adapting this situation to Hall’s example of a magic world, it is possible to have action-at-a-distance in a magic world that is nomologically local, if it is a ‘Law of Magic’ in that world that each spell cast at a some spacetime point P to act at a distance on some point Q also acts on each point along some spatiotemporally continuous path connecting P and Q . The general point is that relations of nomological dependence underdetermine causal structures.

Now, perhaps the kind of causal structures outlined here are too far-fetched to be taken seriously. Or perhaps there is some other way of sharpening up the distinction between situations in which genuine action-at-a-distance is possible and those in which it is not. But if Hall ultimately were forced to rely solely on an intuitive notion of ‘action-at-a-distance-on-the-cheap’ that resists any effort of being rendered more precise, then this would considerably weaken the force of his argument. Nevertheless, in what follows I will grant that there is some such distinction and that it would indeed constitute a considerable cost for a defender of counterfactual accounts, if she had to concede that the kind of cases to be discussed below involve action-at-a-distance.

³ [Frisch 2002] presents an example of this from the physics literature: The so-called *Lorentz-Dirac theory* of classical point charges, is *nomologically* local but, under its standard interpretation, *causally* non-local.

The causal structure that, according to Hall, presents a problem for counterfactual accounts is that of double preventions. Instead of rehearsing Hall's own example—an elaborate science fiction scenario involving World War III fighter pilots—I want to illustrate the structure of double preventions by appealing to a more mundane process involving collisions between billiard balls. The 8-ball is rolling towards the corner pocket and sinks. It would have been prevented from sinking by a collision with the 5-ball, which would have intersected the trajectory of the 8-ball had it not been for a collision between the 2-ball and the 5-ball diverting the trajectory of the 5-ball. Hence the collision between the 2-ball and 5-ball prevents a collision between the 5-ball and the 8-ball, which, if it had occurred, would have prevented the sinking of the 8-ball. Intuitively, the causal structure of this example has the following features. There is one causal 'process' or causal 'chain' of events—the rolling of the 8-ball towards the corner pocket—that takes place unimpeded by its surroundings. And an important reason for the process's being uninterrupted is that a potential interferer or 'threat' to the process—the 5-ball rolling toward a possible collision with the 8-ball—is blocked by a preventing event. The interaction between the billiard balls, which is structurally identical to the science fiction example discussed by Hall, is a case of double prevention.

According to *Dependence*, the collision between the 2-ball and the 5-ball is a cause of the sinking of the 8-ball. For had that collision not occurred, the 8-ball would not have sunk, since it would have been prevented from sinking by its collision with the 5-ball. Yet there appears to be no spatiotemporal continuous causal chain connecting the collision between the 2-ball and the 5-ball and the trajectory of the 8-ball. In the actual world, the 5-ball and the 8-ball never interact. Thus, Hall concludes, *Dependence* and *Locality* conflict. According to *Dependence* there is action-at-a-distance in situations which are physically 'well-behaved' and local, if anything is and, hence, should satisfy *Locality*.

To sum up, Hall's argument can be reconstructed as follows:

- (1) There are double preventions that are governed by local physical laws.
- (2) If a situation is governed by local physical laws, then it also has to satisfy the causal condition *Locality*.
- (3) *Dependence* implies that double preventions violate *Locality*.

(4) Therefore, *Dependence* is false.

Rejecting *Dependence* could take one of two forms. Either we reject *Dependence* outright. Or we restrict its scope to a proper subclass of causal relations and maintain that there are two distinct notions of cause invoked in *Locality* and *Dependence*. According to the causal notion captured by *Dependence*, the first collision in our example is a cause of the sinking of the 8-ball, while according to the causal notion implicit in *Locality* it is not. Hall himself endorses this view, which states that the first collision is a counterfactual cause of the sinking of the 8-ball but does not ‘causally produce’ it.

I will argue that a defender of a counterfactual account need not accept premises (2) or (3). One can reject (3) and argue that, contrary to Hall’s claim, double preventions governed by local laws do not violate *Locality*, since there are spatiotemporally continuous causal chains of events connecting the putative cause and its effect. Perhaps the most obvious variant of this reply appeals to omissions. Applied to our example, one might argue that the collision between the 2-ball and the 5-ball is connected to the trajectory of the 8-ball through a spatiotemporally continuous chain of omissions of the 5-ball and that the omissions are located where 5-ball would have been had it not collided with the 2-ball. In the next section I will argue that there is also a version of this reply that does not rely on reifying omissions as genuine events.

Defenders of *Dependence* can also reject (2). One could appeal to considerations such as the ones discussed above and deny quite generally that a condition of nomic locality implies a condition of causal locality. But I will not pursue this option further here. Alternatively one can grant that there is some connection between nomological and causal conditions of locality, at least in ‘well-behaved’ circumstances, but argue that Hall’s particular proposal for a sufficient condition for action-at-a-distance (*SC*) ought to be rejected in favor of some other causal condition, according to which instances of double preventions governed by local laws are causally local. I will develop this reply in section 4.

3. Lamps and Fields

Consider what happens when I turn off a lamp that is illuminating a screen. Intuitively, my flipping the switch is the cause of the screen’s going dark. This verdict is supported by counterfactual

accounts—according to *Dependence* my flipping the switch is a cause—but also by considerations of blame or responsibility. If the screen goes dark because I flipped the switch, I am responsible for it going dark and am to blame for that outcome. (‘Why did the screen go dark in the middle of the movie?’ ‘It was my fault—I flipped the switch.’) Nevertheless, it may seem that there is no spatiotemporally continuous causal chain connecting my turning the lamp off to the screen’s going dark.

This case is structurally simpler than the double prevention discussed above, for there is no independent causal process corresponding to the motion of the 8-ball that would be prevented from occurring unimpeded, had another prevention not occurred. We can bring out the difference between the two causal structures by considering how they can be represented in terms of a neuron diagram of the kind familiar from the literature on Lewis’s account of causation. The first diagram represents turning the lamp off:⁴

[insert Figure 1 here]

By contrast, a double prevention is represented by the following diagram:

[insert Figure 2 here]

The case of the lamp is a simple prevention characterized by only a single inhibitory link, while double preventions have two such links. Let us note, then, that there are causal structures simpler than the one considered by Hall for which the problem to which he points appears to arise with equal force: even simple preventions involve non-local causation, according to *Dependence*, or so it may seem.

Hall’s ‘favorite failed strategy’ of trying to reconcile *Dependence* and *Locality* is to appeal to chains of omissions. He devotes the most space to discussing and rejecting this strategy in [Hall 2002] and it is the only alternative he discusses in [Hall 2004]. Applying this strategy to the example of the lamp, one might try to avoid the putative conflict between *Dependence* and *Locality* by arguing

⁴ The two dark neurons are firing, arrows represent positive causal links and the connection with dot represents an inhibitory link. The lower neuron represents the flipping of the switch. While the voltage at the outlet is non-zero (represented by the top left neuron firing), the flipping of the switch results in the lamp being off (represented by the middle neuron on the top, which is not firing) and the screen being dark (represented by the neuron to the right).

that the event of turning off the switch and the event of the screen's going dark are connected by a spatiotemporally continuous causal chain of omissions of photons or light pulses. Similarly, in the case of double preventions discussed above one could point out that in *counterfactual* worlds in which the collision between the 2-ball and the 5-ball does not occur the trajectory of the 5-ball would have intersected that of the 8-ball. And one might then argue that, corresponding to the trajectory of the 5-ball in the counterfactual world, there exists a chain of absences or omissions of the ball in the actual world that connects the trajectory of the 8-ball with the collision the 2- and 5-balls.

A major part of [Hall 2002] is devoted to rebutting this strategy. I am not convinced that Hall's extensive and subtle discussion of the omissions strategy really does succeed in showing that the strategy fails. Nevertheless there are serious worries about the ontological status of omissions and it would be preferable if a defender of *Dependence* would not have to reify omissions to meet Hall's challenge. Fortunately a different strategy is available, for contrary to our initial impression there *is* a local causal structure in the case of the lamp, about which talk of omissions of waves or photons provides information.⁵

We do not think that the lamp acts at a distance on the screen, but rather that any effect that the lamp might have on the state of the screen is mediated by an electromagnetic field. But in the context of a field theory the state of the field is thought of as a property of spacetime quite independently of whether the field happens to be zero at a given spacetime point or have some non-zero value. The set of basic events, in a field-theoretic framework, is the set of instantiations of properties at spacetime points: events are constituted by fields having certain values at specific points. For a field to be zero at a particular spacetime point is a genuine event, just as it is for the field to have some non-zero value. As I turn off the lamp, a region where the electromagnetic field is zero spherically diverges from the lamp. Just as a non-zero field pulse diverging from the lamp as I turn it *on* constitutes a localized causal chain linking lamp and screen, the propagation of the region with zero fields constitutes a continuous causal chain whose links consist of the instantiation of spatiotemporally localized properties.

⁵ [Beebe 2004] argues that we construe talk of omissions as providing us with information about causal structures that are comprised of non-omission events.

There is an interpretation of classical electrodynamics, due to John Wheeler and Richard Feynman, that takes the electromagnetic field to be a mere calculational device and, hence, treats the interaction between lamp and screen as an interaction-at-a-distance. If we assumed such an action-at-a-distance electrodynamics—that is we assumed that the interaction between lamp and screen was not mediated by an electromagnetic field—the fact that my turning the lamp *off* is not linked to the appearance of a dark spot on the screen through some spatiotemporally continuous causal process would not seem to be especially troublesome—it would be no more troublesome than the appearance of a bright spot immediately after I turn the lamp *on*. If the physics is non-local, then we would expect causal relations to be non-local as well. Indeed, it is often argued (in contexts where the status of omissions play no role) that it is one of the main virtues of field theories that they allow us to rescue locality (in various senses). It seems that Hall's worry arises only if, on the one hand, we assume, as we normally do, that the underlying physics *is* local due to the existence of electromagnetic fields mediating electromagnetic influences, but, on the other hand, we do not realize that absences of field excitations also correspond to localized properties of spacetime that propagate locally *in the very same way* as do non-zero field excitations. Thus, if we do not treat electromagnetic interactions field-theoretically, there is no conflict between *Dependence* and *Locality*, because we do not expect *Locality* to be satisfied. But neither is there a conflict if we adopt a field theory, for we then have a framework that allows for the local propagations of regions with zero fields just as it does for non-zero field excitations.

At this point it seems that, contrary to what our initial intuitions might have suggested, the lamp-and-screen-example is not an instance of a failure of *Locality*, precisely because our current physics suggests that the case be treated field-theoretically. But this, it may seem, leaves the billiard ball case I introduced above untouched. Yet modern spacetime theories extend the field-theoretic framework to cover so-called 'matter fields' as well. Thus, taking a cue from present day fundamental physics, we could generalize our field-theoretic treatment to the billiard ball case and others like it by modeling the trajectories of material objects in terms of the time-evolution of the distribution of matter fields and treating talk of omissions as providing information about the instantiations of properties of spacetime. If we start with spacetime points as our basic ontology, then

the absence of a certain kind of object in a certain region corresponds to a genuine event—a field having zero value in that region—and this constitutes a legitimate property that ought to be allowed to play a causal role just as non-zero field values do. That is, in the billiard ball case, the collision between the 2-ball and the 5-ball comes out as a local cause of the sinking of the 8-ball, because the collision is linked to the trajectory of the 8-ball through a spatiotemporally continuous chain of events that consist in certain spatiotemporal regions being relatively devoid of matter (or the ‘billiard-ball-field’ being zero). While this way of conceptualizing the situation may seem odd initially, this is precisely how the presence or absence of matter in spacetime regions is thought of in General Relativity.

But is it really the case that on a counterfactual account there is a spatiotemporally continuous causal chain linking the collision of the 2-ball and the 5-ball with the sinking of the 8-ball?⁶ One might worry, for example, that a counterfactual approach cannot deliver the verdict that the ‘boring’ event of a region R adjacent to the 8-ball being relatively devoid of matter is a cause of the 8-ball’s sinking, for it might not be clear what state of R is supposed to take the place of the boring state in counterfactual worlds closest to the actual world. But, first, this worry is a problem only for certain variants of counterfactual accounts—chiefly among them David Lewis’s original account—and, second, where the worry does arise it constitutes a general problem for the accounts at issue and arises independently of considerations of locality.

According to Lewis’s account, the sinking of the 8-ball causally depends on the event that region R is in a ‘boring’ state exactly if it is the case that there are no worlds in which the ‘boring’ event does not occur in R and in which the 8-ball does sink that are closer to the actual world than any world in which the boring event does not occur and the 8-ball does not sink. And it seems doubtful that this could shown to be the case. In particular, it is doubtful that we could establish that the event replacing the boring event in the closest possible worlds is the presence of the 5-ball with the right momentum to knock the 8-ball off its trajectory.

⁶ This objection is due to an anonymous referee for this journal.

Notice, however, that the problem as to what counterfactual event is to take the place of an actual event is one that Lewis's account faces much more generally, as Hall himself has argued [Collins et. al. 2004: 19-21]. Take the counterfactual 'If the 2-ball had not collided with the 5-ball, the 5-ball would have collided with the 8-ball.' What takes the place of the collision between the 2-ball and the 5-ball in the worlds closest to the actual world? Would there be another ball with which the 5-ball collided? Would a player reach in with her hand? Would the region, where the 2-ball was located, contain a vacuum? Lewis has suggested that the counterfactual situation we are to imagine is one where the actual event 'is completely and cleanly excised from history, leaving behind no fragment or approximation of itself' [Lewis 2004: 90]. But it is not clear to what extent these remarks offer a solution to the problem, for each excised event has to be replaced by some *other* counterfactual event, and it is not clear that Lewis's similarity metric determines uniquely what that event ought to be. The problem, then, as to what counterfactual events to consider in evaluating Lewisian counterfactuals is general and arises for local structures as well.

Moreover, there are alternatives to Lewis's original counterfactual account of causation that do not face this problem. One strategy is to propose that e counterfactually (and hence causally) depends on c , just in case it is true that if *some* other event c' (in a set C) had occurred instead of c , then e would not have occurred. That is, in probing for counterfactual dependence, we focus on a whole range of alternatives to c and ask if any one of them would result in e 's not occurring. Lewis's recent account of causation as influence [Lewis 2004] adopts this strategy, as does James Woodward's interventionist account of causation [Woodward 2003]. On Woodward's account, the appropriate counterfactuals for evaluating causal claims are those that have to do with the outcomes of hypothetical interventions. A sufficient condition for causation, according to Woodward, is the following (where the causal relata are the values of variables rather than event):

If (i) there is a possible intervention that changes the value of X such that (ii) carrying out this intervention (and no other interventions) will change the value of Y [...], then X causes Y . [Woodward 2003: 45]

Thus, on Woodward's account, if a counterfactual of the form 'If an intervention set the value of X to x^* , then the value of Y would be y^* ' is true, where x^* and y^* are non-actual values, then X is a cause

of Y . Applied to our example, the ‘boring’ event adjacent to the trajectory of the 8-ball comes out as a cause of the 8-ball’s sinking, since there are possible interventions which would result in the 8-ball’s not sinking—in particular an intervention that resulted in the presence of the 5-ball with the right momentum to knock the 8-ball off its trajectory. Notice that Woodward’s account is still committed to *Dependence*, since the value of Y depends counterfactually on the value of X exactly if there are true counterfactuals of the form above, and, thus, counterfactual dependence, on this account, implies causal dependence. But Woodward is not committed to the idea that counterfactuals of the form ‘If the value of X had not been x , then the value of Y would not have been Y ’ have context-independent truth conditions.

In this section I have presented one strategy that a defender of a counterfactual account can employ in responding to Hall’s challenge. If two events are causally related according to *Dependence* and the underlying physics is local, then Hall’s condition for action-at-a-distance *SC* is not satisfied, since there is a spatiotemporally local chain of events: the instantiation of field values at spacetime points connecting the two causally related events. I will now turn to a second strategy for meeting Hall’s challenge.

4. A counterfactual criterion of causal non-locality

The second strategy that I want to discuss consists in proposing a rival condition of non-locality, which is not in tension with *Dependence*. A sufficient condition for causal action-at-a-distance, I submit, is the following:

(AD) If an event c lies in the causal past of an event e and there is a spacetime shell S surrounding e but not c such that the occurrence of e counterfactually depends on the occurrence of c even when we keep the physical state of S fixed (except for events on S that lie in the causal future of e), then c is a cause acting at a distance on e .

AD is satisfied neither for cases of double prevention nor for cases of simple preventions. If we keep fixed that the 5-ball is absent from some spacetime shell surrounding the 8-ball, then counterfactually

varying whether the 5-ball collides with the 2-ball makes no difference to the trajectory of the 8-ball. Similarly, if we keep fixed that there is an electromagnetic wave traveling toward the screen, then varying whether or not I turn the lamp off has no effect on the fact that the screen is illuminated. Both cases are causally local according to *AD*.

AD is most naturally seen as part of an interventionist account of causation, such as the one developed in [Woodward 2003], which take counterfactual interventions into a system from the outside to track the causal relations among the system's components. One might worry that if events on *S* are in the causal future of *c*, then changing whether *c* occurs without affecting the state of *S* may be nomically impossible: it is nomically impossible, it seems, for there to be an electromagnetic wave traveling toward the screen if the lamp is turned off. An interventionist account can help to ease that worry. For it might be possible to postulate multiple interventions, some of which are acting on the state of *S* directly and thereby allow us to 'freeze in' that state while varying the occurrence of *c*. Moreover, if it is possible to isolate distinct causal routes, even hypothetical interventions that are not nomically possible need not be problematic. As long as it is possible to restrict the violation of the laws to the link between *c* and the state of *S*, such that any direct link between *c* and *e* or between *S* and *e* is still governed by the relevant laws, it should still be possible meaningfully to apply *AD*.

Corresponding to the pair of conditions of nomological locality NL_1 and NL_2 , we can also introduce a second counterfactual condition of causal locality—a finite propagation condition:

(*FP*) An event *e*, occurring at time *t* and location *x*, has no non-local causes exactly if at all times $t-\Delta$ there is a complete set of causes of *e* occurring in a finite neighborhood δ of *x*, for positive δ and Δ , and the size of δ goes to zero as Δ goes to zero.

FP disagrees with *AD* in the case of causes that propagate infinitely fast but not across spatial gaps, as for example in the case of forces acting on a rigid body. According to *FP*, rigid body mechanics is causally non-local. Since forces propagate infinitely fast through a rigid body, the region within which a complete set of causes of the motion of one end of a body is located never becomes smaller than the size of the entire body even as Δ becomes infinitesimal. By contrast, the theory is causally local according to *AD*. Each spacetime shell around a point of a body includes 'links' of all causal

chains that propagate through the body. Unlike in the case of *FP*, the spacetime shells in *AD* also include events that occur simultaneous with *e*, and hence they ‘pick up’ causal processes propagating infinitely fast towards *e*, as long as they do not jump across spatiotemporal gaps.

Hence, *FP* is satisfied exactly if causes are local in *both* of the following two distinct senses: Causal influences propagate at a finite speed; and there is no causation across spatiotemporal gaps. *AD*, by contrast, is only concerned with the second of these constraints and is silent on the speed of causal propagation.

Now, Hall thinks that there is an extremely quick and decisive counterexample to counterfactual conditions of non-locality such as *AD*. Imagine that a magician is just about to cast a spell that would act at a distance but in the actual world is prevented from doing so, since I wrest his wand from his hand. Then, Hall maintains, there is no action-at-a-distance in the actual world, even though there would have been action-at-a-distance in the closest or most relevant counterfactual world. But according to a counterfactual criterion such as *AD* there *is* action-at-a-distance in the actual world. Hence the criterion fails—it conflates ‘cases where action at a distance *would have* occurred with cases where action at a distance *does* occur.’ [Hall 2002, 293]

Hall’s argument relies on the assumption that the fact that there *would* be action-at-a-distance had a certain event not occurred does not imply that there *is* action-at-a-distance in the actual world. If the magician had cast his spell, there would be action-at-a-distance. But this does not imply that my wresting the wand from the magicians hand also is an event that acts at a distance. Hall’s assumption, one might think, captures an intuitively obvious distinction. Thus, by showing that a counterfactual account is blind to the distinction between full-blown action-at-a-distance, when the spell is actually cast, and the ‘action-at-a-distance lite’ that merely consists in preventing the spell, Hall has at the very least succeeded in showing that counterfactual accounts come at a cost.⁷ But I want to argue that the cost is minimal or even non-existent.

Generalizing Hall’s assumption to all causal relations—both local and non-local—the assumption becomes this: the fact that causation *would* occur does not imply that causation *does*

⁷ This point has been pressed by an anonymous referee for this journal.

occur. That is, according to the assumption there are triplets of events c , d , and e , such that: c and d occur in the actual world; d would cause e if c did not occur; and neither c nor d is causally related to whatever actual event e' occurs where- and whenever e would occur if c had not occurred. How do we evaluate the truth of this assumption? It is not clear that intuitions unequivocally support it.

Either d is also a cause of e in the actual world, or it is not. If it is, then there is causation in the actual world; but if it is not, then it must be that c 's occurrence prevents d from causing e and then it might seem plausible that c is causally related to whatever event e' occurs in the actual world instead of e .

We can ask whether particular accounts of causation imply the assumption. A *counterfactual version* of the general assumption is indeed true—counterfactual dependence in counterfactual worlds does not imply counterfactual dependence in the actual world—but it is true for an interesting reason: there is no counterfactual dependence in the actual world even when there is counterfactual dependence in non- c worlds in cases of overdetermination. And such cases pose a problem for the necessity of a counterfactual criterion and do so precisely because we think that overdetermination involves causation even though there is no counterfactual dependence. The assumption is also true according to *process accounts of causation*, since preventions of causal processes need not themselves involve a causal process. But in this case, too, the verdict might be thought to present a problem for that account of causation.

Hall's case of the thwarted magic spell is, of course, a case of prevention. What, then, is the correct causal analysis of preventions? Hall maintains that cases of prevented action-at-a-distance should not themselves count as involving action-at-a-distance. But when we compare Hall's non-local case with preventions that do not involve potential action-at-a-distance it is far from obvious that we ought to share Hall's intuitions. Imagine, for example, that I am about to grab my umbrella upon leaving my office, but that you play a practical joke on me and hide the umbrella before I am able to fetch it. I leave without umbrella and get wet in the rain. Had you not hid my umbrella, I would have taken it with me and I would have opened it as I stepped outside. According to *Dependence*, your hiding the umbrella is a contributing cause to my not opening the umbrella and, thus, to my getting wet. And I think that there are fairly strong intuitions that agree with this verdict. ('It is your fault that I got wet,' I might complain, or even 'Your prank caused me to get wet.')

As a second case consider once again the example of my turning off a lamp, which we discussed above. According to *Dependence*, my flipping of the switch is a cause of the screen going dark. And, as I said above, this verdict seems to be supported by our intuitions. ('Look at all the troubles you have caused,' you might complain to me after my flipping the switch interrupts the movie and mayhem breaks out in the audience.) These two cases are local, according to *AD*. Yet aside from that fact all three cases—the two local examples and Hall's case of the magician—have the same causal structure, represented by the first neuron diagram above.

Hall thinks that in the case of the magician, since no spell was actually cast, the preventing event does not act at a distance and, hence, simply is *not a cause* of any events further along the putative causal chain. 'There is a difference, after all,' he says, 'between casting a spell and merely failing to cast one.' [Hall 2002: 293] Yet if Hall treats equivalent causal structures on an equal footing, then he is committed to denying in general that preventing events can act as causes: my flipping the switch, he would have to maintain, is not a cause of the screen's going dark, since no non-zero electromagnetic wave was actually emitted ('there is a difference, after all, between emitting a non-zero electromagnetic wave and merely failing to emit one'); your hiding my umbrella is no cause of my getting wet, because I did not actually take the umbrella along ('there is a difference between taking along an umbrella and merely failing to take one', and indeed there is: the difference may involve one's getting wet); and, finally, in the neuron diagram the inhibitor neuron's firing is not a cause of the last neuron to the right not firing.

Hall's intuition that preventions do not involve causation (and that to assume otherwise involves a confusion between cases where causation *would* occur and cases where causation *does* occur) seems in fact to be motivated by a process or conserved quantity account of causation. Yet it is not obvious why a defender of *Dependence* should feel compelled to share this intuition. According to counterfactual accounts of causation, *there is nothing more to causal dependence than counterfactual dependence*. Thus, events represented in a neuron diagram as inhibitors are causes just as much as those represented as activators. Prima facie it seems that a counterfactual criterion gets the causal structure exactly right here. Indeed, process accounts are often criticized for not allowing preventions do involve causation and even some advocates of process accounts recognize that there

are strong causal intuitions *in favor of* treating preventions as causes (and, thus, appeal to notions such as that of *quasi-causation* to characterize preventions (see, e.g., [Dowe 2001: 221]).

Moreover, in order to establish the thesis of the present paper—that Hall’s considerations of locality do not present an *independent* challenge to counterfactual accounts—I do not have to show that *Dependence* gets the causal structure right in cases of preventions. It is enough for my purposes here to establish that there are causally *local* cases with the same causal structure as the example of the thwarted magic spell. If there were good reasons for denying that the firing of an inhibitor neuron is a cause of the non-firing of neurons to its right, then these reasons would pose a problem for *Dependence*. But this would represent a *general problem for Dependence* that arises quite independently of considerations of non-locality.

So far I have assumed that Hall’s claim that the thwarted magic spell does not involve causation can be extended to local cases of prevention. Yet Hall might maintain that there is a crucial distinction between preventions of locally acting causes and preventions of action-at-a-distance: despite the equivalence in causal structure the former involve causation while the latter do not. But defenders of counterfactual accounts can and should reject such a distinction as being *ad hoc*, especially in light of the fact that Hall is not denying the possibility of action-at-a-distance in general. If Hall wanted to draw such a distinction, he would have to maintain that *inhibitors* are causes only when they are locally connected to their putative effects while allowing that *activators* can act as causes both locally and at a distance. That is, he would be forced to maintain that the correct causal interpretation of a given neuron diagram depends not only on which of its interactions involve inhibitory or stimulating signals, but also on whether an arrow to the right of an inhibitory signal represents a local interaction: Inhibitory signals can be causes of the state of a neuron to the right of such an arrow, but only if the arrow represents a local interaction; while activating neurons are causes, independently of whether the interaction is local. It is difficult to see how such a distinction could be motivated.

6. Conclusion

Hall claims that there is a conflict between counterfactual accounts of causation and a causal condition of locality. If counterfactual dependence implies causation, he argues, then non-locality can be had on the cheap. He takes this to suggest that we should distinguish between two causal relations—a counterfactual relation that is simply not concerned with questions of locality and a relation of causation as production, which has the property that only ‘genuinely’ non-local processes come out as causally non-local. I have discussed two strategies for responding to Hall’s challenge. First, one can adopt a framework in which the fundamental events are given by suitable fields having specific values in regions of spacetime. Second, I have proposed a counterfactual criterion as sufficient condition for causal non-locality that draws the distinction between local and non-local causation in intuitively the right place.

What is the relation between these two strategies? In the context of a field-theoretic framework *SC* implies *AD*. If there is a cause that is not connected to one of its effects through a spatiotemporally continuous causal chain then the occurrence of the effect varies counterfactually with that of the cause, even if we keep the state of some spacetime shell surrounding the effect but not the cause fixed. But *AD* does not imply *SC*. In cases where a cause is connected to one of its effects along two distinct causal routes exactly one of which involves action-at-a-distance, *AD* implies that the cause acts at a distance, yet *SC* is not satisfied. Yet there is a condition in the neighborhood of *SC* that renders the different approaches equivalent and even provides necessary and sufficient conditions for causal action-at-a-distance (within the context of the field-theory strategy):

(*SC**) We have a case of action-at-a-distance exactly if we have a cause, which has at least one effect that is connected to it via at least one causal chain that has no spatiotemporally continuous completion.

Here a causal chain *C* is any sequence of events such that any event in the sequence is an effect of the event that immediately precedes it and a completion *O* of a causal chain *C* is itself a causal chain such that we can obtain *C* from *O* by deleting some of *O*’s members.

Despite the fact that the two strategies I discussed agree in their verdicts, I prefer the second strategy, which appeals to *AD*. It has the advantage that it requires fewer ontological commitments.

Once we accept *AD*, it does not matter whether we also adopt a field-theoretic framework for our event ontology.

But why, one might ask, should the issue of locality hinge on the state of a shell *S* surrounding an event, even though this shell might contain no events that are causally related to the effect in question? The answer is that even if we are not willing to take regions surrounding an effect to be populated by ‘zero-field events’, the claim that there are no causally efficacious events occurring in this region does provide us with information about causally salient features of the situation. If we keep the state of a shell surrounding an event fixed, then we keep fixed both what events are and what events are *not* occurring on the shell. According to *Dependence*, turning off the lamp is a cause of the screen’s going dark, and according to *AD* it is a locally acting cause, since, according to the laws, the only way in which my flipping the switch can have an effect on the screen is by having an effect on what events occur on shells surrounding the screen. Similarly, the sinking of the 8-ball causally depends locally on the collision between the 5- and the 2-ball, since the only way in which the motion of the 2-ball can have an effect on the sinking of the 8-ball is by having an effect on events occurring in the neighborhood of the 8-ball.

Hall’s locality condition seems to be motivated by thinking about causation in terms of causal processes during the course of which events ‘produce’ or ‘bring about’ their effects. Yet defenders of counterfactual accounts of causation wish to deny that causation has this character. Thus, it seems to me that what Hall’s paper points out is not that there is a tension between considerations of locality and *Dependence*, but rather that there may be a tension between a *process* criterion of locality and a *counterfactual* condition of causation. But there is no reason why the claim that there is a tension between a *non-counterfactual* locality condition and *Dependence* should be of concern to a defender of *Dependence*. For an advocate of a counterfactual account of causation can simply reject Hall’s locality condition and insist that in order to convince her of a tension between locality and *Dependence*, Hall would have to show that there is a tension between *Dependence* and the kind of causal locality criterion which she would accept—that is, a *counterfactual* criterion such as *AD*. Yet it is difficult to see how there could be such a tension. Thus, at least as far as considerations of

locality are concerned, there is no need to multiply causal relations—a counterfactual condition alone is all we need.

University of Maryland, College Park

REFERENCES

- Beebe, Helen 2004. Causing and Nothingness, in *Causation and Counterfactuals*, ed. John Collins, Ned Hall and L. A. Paul, Cambridge, Mass.: MIT Press: 291-308
- Belot, Gordon 1998. Understanding Electromagnetism, *British Journal for the Philosophy of Science* 49/4: 531-55.
- Cohen, I. B. 1978. *Isaac Newton's Papers and Letters on Natural Philosophy*, Cambridge, MA: Harvard University Press.
- Collins, John, Ned Hall and L.A. Paul 2004. Counterfactuals and Causation: History, Problems, and Prospects, in *Causation and Counterfactuals*, ed. John Collins, Ned Hall and L. A. Paul, Cambridge: MIT Press: 1-58
- Dowe, Philip 2001. A Counterfactual Theory of Prevention and 'Causation' by Omission, *Australasian Journal of Philosophy*, 79/2: 216-26
- Frisch, Mathias 2002. Non-Locality in Classical Electrodynamics, *British Journal for the Philosophy of Science* 53/1: 1-19.
- 2005. *Inconsistency, Asymmetry, and Non-Locality*, New York: Oxford University Press.
- Hall, Ned 2002. Non-locality on the Cheap? A New Problem for Counterfactual Analyses of Causation, *Nous* 36/2: 276-294.
- 2004. Two Concepts of Causation, in *Causation and Counterfactuals*, edited by John Collins, Ned Hall and L. A. Paul, Cambridge: MIT Press: 225-276.
- Lange, Marc 2002. *An Introduction to the Philosophy of Physics: Locality, Fields, Energy, and Mass*, Oxford: Blackwell.
- Lewis, David 2004. Causation as Influence, in *Causation and Counterfactuals*, edited by John Collins, Ned Hall and L. A. Paul. Cambridge: MIT Press: 277-290.
- Woodward, James 2003. *Making Things Happen: A Theory of Causal Explanation*, Oxford: Oxford University Press.